ICANN71 | Virtual Policy Forum – RSSAC Work Session 1
Tuesday, June 15, 2021 – 16:30 to 17:30 CEST

OZAN SAHIN: Hello, and welcome to the RSSAC Work Session. My name is Ozan Sahin, and I am the remote participation manager for this session. Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior.

During this session, questions or comments submitted in the chat will only be read aloud if put in the proper form as noted in the chat. I will read questions and comments aloud during the time set by the chair or moderator of this session. If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly and at a reasonable pace. Mute your microphone when you're done speaking.

This session includes automated real-time transcription. Please note that this transcript is not official or authoritative. To view the real-time transcription, click on the Closed Caption button in the Zoom toolbar.

With that, I will hand over the floor over to Ken Renard.

KEN RENARD: Thank you, Ozan. Thanks, everyone, for joining today. We have some new folks in the participation in the room here today, so I'll give a brief

background of where we are, what we're doing. Then we can dive into the work that we're scheduled to do today for this work party.

So, the work party here is the Tool on the Local Perspective. It's meant to define a tool to measure the Root Server System from a particular vantage point on the Internet. It is complimentary to the RSSAC047 work which is measuring from well-connected areas. This is meant for somebody anywhere on the Internet to see what the Root Server System looks like from their exact vantage point.

We've been working on this work party, in this document, for a little over a year now, and we're really at the point where we're wrapping it up. We basically have one section and an associated subsection in the document to work out today, as well as some editorial things to go through that have been brought up in the past couple of days.

With that, I will walk briefly through the document and describe what we're doing. And then we can dive right into the discussion on the repository.

So, again, this is the Tool for Local Perspective. We have, in Section 2, a couple of use cases. Who are the people that would use this tool? The first use case is informing the determination of an underserved area. So, we use this tool to measure how the Root Server System looks in order to eventually determine if an area is underserved and maybe worthy of placing a new root server instance in that area or close by.

The tool only provides data into that decision. It's not making the decision. It's not the case where if the measurement is greater than X,

then it automatically gets a new instance. It's only to inform that decision.

A similar but different perspective on that is evaluating third-party requests for hosting an AnyCast instance. So, as Root Server Operators get requests to host a new instance of their root server, they may use this tool to get some measurements and determine how good or bad the Root Server System is from that particular location and then determine whether it's worthwhile to place an instance there.

The third use case we describe is a recursive operator that may want to use this tool to look, maybe, at various parts within their own network at how the Root Server System looks and see if they want to change any of their configuration—maybe it's routing, maybe it's additional load or additional recursive resolvers in their network— somehow to best take advantage of the Root Server System and the measurements from their particular vantage points.

So, we describe those use cases, and then we go through and talk about measurements which can support informing those use cases. We look at a couple things here. The tool will collect a timestamp when this is run. But then we get into DNS query latency. We have a couple different types of measurements that we use to look at different queries recording the timing.

We also look at routing and conductivity, usually doing a trace route to see what paths might be. And those can support decisions based on if you're getting the optimal route. And we also look at the measurement source. So, we want to get an IP address of the measurement location.

So, using DNS to sort of, essentially, abstract out any NAT or effects of private IP address space that would confuse that. So, we try and get where this measurement was run from.

So, the idea is that we can gather this information and then push this data to a data repository where it can be analyzed, where it can be compared. For example, in the use case where we're looking at underserved areas or trying to inform that decision, we can compare to what we think are well-connected sites versus maybe not well-connected sites and look at those numbers. And, given enough data, we can really determine, "Hey, this is a particularly bad location with respect to performance of the Root Server System."

So with that, if there are any questions, feel free to raise a hand. That's kind of the background of the document, and the current topic of discussion is the Data Repository and Recommendations.

So, Paul Hoffman had made several good comments in here. I'd like to address those, but I think the best use of our time is to dive into the discussion of the data repository and then go to those recommendations from Paul. I think they're all very good ones and we can mostly just accept those.

Andrew.

ANDREW MCCONACHIE: Hey, Ken. So, I had a question about the highlighted "should." "All users of the tool should publish their results to the repository." You saw my comment, and it reads a bit to me like users are expected to or must. I

mean, I know we're not using RFC language here, but do we want to soften that? Or what is the requirement we're going to place on users who run the tool to send the results into the data repository?

KEN RENARD: Okay. And which section is that in? Is that in the Repository …?

ANDREW MCCONACHIE: Yeah. That's in Section 5.

KEN RENARD: Okay.

ANDREW MCCONACHIE: It's “All users of the tool should …”

KEN RENARD: Yes. So, I was actually trying to use RFC language there—“should” not being an absolute requirement. But I think we really want absolutely as much as possible for people to publish their results for whatever privacy reasons. I don't think that we should absolutely require it. So, by having it as an opt-out of publishing approach, I think that would get us as much data as possible but still give the user some level of privacy and still be able to use this tool.

Paul.

PAUL HOFFMAN: I have the same concern Andrew does, especially given the three use cases in Section 2. That is that some of the people who are expected to be using the tool are using it just for their own "I'm a resolver or operator and I just want to see what's going on around me."

If there was any thought that my just running the tool was going to cause data publications somewhere else, I would be very hesitant to use the tool. I think the wording that you have in this section is aimed at researchers, which I think is reasonable to ask the research community to, "Once you have measured things, please publish them. And here's an easy way to do it all."

But if I am a resolver or operator or, quite frankly, if I'm in RSO, I wouldn't want somebody to say, "And here's how you're going to use the output. Here's how the output of what you're running is going to be used in the general world." Thank you.

KEN RENARD Okay. Thanks, Paul. Andrew, did you have more?

ANDREW MCCONACHIE: I did. I'm going to make a suggestion in the document really quick. Just because these are requirements for the document—or for the tool itself— not so much for the user. But the tool should encourage users to publish their results to the repository. That way it's kind of like, lower down, we talk about opting out. So, here we can kind of say that the tool is setting a default. So, it gives users some …

I think it should satisfy your concern that we want, as much as possible, users to contribute their data, but they can still opt out if they wish.

KEN RENARD: Yes, that does address what I'm thinking. We'll never require people to publish. It's just a matter of how bad do we want the data … Even making decisions on underserved requires a good bit of data, so it should be strongly encouraged. I'm fine with what you've put in there.

Wes.

WES HARDAKER: Yeah. Just a point of reference. I implemented a DNSSEC checking tool a long time ago that I published on Android and desktops and things like that, that allowed people to upload the results of their studies. And I actually received quite a bit of data from that. And the way I did it was, you run the check. It gives you the results. And at the bottom is a little button saying, "Submit these results."

And I think the average person probably will do it. And at least I have decent experience that you can collect a fair amount of data that way. It just needs to be done properly with a UI system that sort of encourages it.

KEN RENARD: Okay. That's good news. You could get three free root server queries if you publish some sort of incentive. All right. That sounds good.

Wes, you had also talked about your comment in the document about an explicit goal. Yeah, I know we've shied away before from building this as a research tool. We have the use cases. We used to have research as one of the use cases, but we took that out.

WES HARDAKER: Yeah. I wasn't trying, necessarily, to say research. I mean, if users of the tool want to compare their results to other people nearby or to get a bigger picture to help them understand their local network, it's not likely going to be just the results they want to see. So, having the repository in the system as open as possible allows better use. And, yes, it allows for researchers to use it, too, even if that's not an explicit goal. But I think my point was more that it should be as open as possible unless we have a reason that it shouldn't be.

KEN RENARD: Okay, yes. Saying that it will be used by researchers, but used by … If I'm running this to understand how the RSS looks at my site, yeah, I will want to see how it compares to others.

So, how do we describe this as being available to researchers as well as others? My reply to your comment was along the lines of, how do we determine who's a real researcher and who's just trolling around. I think just stating that it's open avoids that decision. Or stating it should be as open as possible.

WES HARDAKER: Right. So, I don't know that we need to specify … I don't think we need to build a use case if we believe we have multiple use cases, and I think that we do. Right? I don't think we need to call out, "We're doing this because of researchers." Right?

We want to say that, in the interest of the Internet good, things should be generally open except for private, sensitive information that may be particular to individual data submissions or users. And this may benefit people studying their local regions as well as researchers studying the global RSS as a whole to whatever. Right?

I don't think we need to call out that we're doing it because we like these types of people.

KEN RENARD: Agreed. Yeah, the conflict in my mind right now is between as much use for the good of the Internet as possible versus somebody being scared to publish their data because they don't know who's going to look at it. But your experience with your tool, Wes, has shown that people are actually likely to share.

Paul.

PAUL HOFFMAN: So, two things here. One is that I very much agree with Wes and the way that his tool works, which is that one doesn't have to opt out. One ops in but is encouraged to opt in, and may be encouraged in such a way that says, "And here is why this might be useful."

But the other is that, given the three audiences above—which, by the way, doesn't actually include just plain old researchers—I think it would be sad if this tool is finished and published and then someone complains and says, "I just ran it to see how the tool works, and all of a sudden, I see my data in this repository and I don't see a way of taking my data out. And they didn't tell me really clearly," and such like that.

So, I think opt-in gives the tool a better sort of marketing value. That is, "We aren't going to take your data without knowing. But, boy, if you want to help the world, it's easy to do so." Thank you.

KEN RENARD: Thanks, Paul. Okay. So, it sounds like we want to go to a very strongly encouraged opt-in versus an opt-out, stating that the data should be open to researchers and to other users of the tool, potentially just an aggregate.

All right. Trying to go through other comments in this area.

FRED BAKER: One additional thought in that thread. When you say "researcher," I think Ph.D. candidate. I think what you mean is someone doing research, and it might be worthwhile to consider how to say that.

KEN RENARD: Okay. "One doing research." Okay, Wes, you had the comment about structured and unstructured data.

BRAD VERD: Hey, Ken. Going back to the previous thread, also, I guess I'm going to … For the record, I'll just politely disagree with maybe everyone. I'm not sure. I hear all the conversations being made and it's not that I disagree with them in the spirit, but I guess what I'm—[inaudible], I should say, is available via—

WES HARDAKER: Brad, you cut out. Can you restart since you disagree with everybody—after that?

BRAD VERD: Can you hear me? Sorry. I'm trying to …

WES HARDAKER: Yeah. It's okay. You just cut out for about a good 5-10 seconds.

BRAD VERD: That's the joy of the Internet. I'm trying to say that most of the functionality that we've talked through all of this, we have all agreed, is in RIPE Atlas where the data is available. And we've also stated that what we're trying to accomplish here is for use not necessarily by researchers, though we continue to talk about researchers. The primary uses for the RSOs to have data to make informed decisions about locations for AnyCast instances.

And with that in mind, I feel that the more data you have, the more informed the RSOs are to make decisions. I can't even begin to enumerate the amount of applications that I download, which is essentially what we're proposing here—some tool that either you run or download and whatnot—that are all, you have to opt out if you don't want your data shared.

Because you're getting the service, meaning the application, for free, so to speak, because the engineering work has been done for you. The obvious ones to point out are like search and social media and whatnot. All that data is just used and you, in some cases, can't even opt out. So, I guess I'm just … The conversation is interesting to me, but in the spirit of …

What I thought we were after here was to get RSOs to make informed decisions on new instances to better serve the globe. And I'll leave it at that. Thanks.

KEN RENARD: Thanks, Brad. Looking at our use cases, the first two—informing in underserved areas, evaluating third-party requests—yes, those are directly related to supporting RSO decisions. The recursive operator use case is one where it's just [a user or a] site on the Internet, trying to understand their perspective. And that's the one where they're more likely to make a decision not to publish because it would no longer inform the goal.

I'm trying to remember the order here. I think Paul Hoffman was next.

PAUL HOFFMAN: Thank you. So, I can agree—I can't. I agree with Brad. If the purpose of this work, or the main purpose is to inform the RSOs, then automatic data collection could be part of this. If that's the case, though, then I think you almost need to drop the third use case of resolvers just wanting to check themselves.

And I know this goes back to the very beginning in the Statement of Work for this work, which is, "Oh, we can also do this and we can also do this." But opt-in versus opt-out is sort of inherent in this. My suggestion of making it opt-in with encouragement was pretty much based around the third use case, the resolver or someone who thinks that they're a resolver or an operator—they don't even know what that really means—running the software and then being concerned.

If, in fact, the primary use cases are for the better understanding by RSOs, then I think it's fine to have it automatically send data as long as that is said up at the top. When you install the program and you first run it, it says, "This is what the program is going to do. It's going to collect this data, and it will send it here so that other people can see it," and such. And then they click Run or Cancel.

I think that's just fine. But in that case, I think you really need to reword the third use case to sort of downplay it. Thank you.

KEN RENARD: Thanks, Paul. I thin Wes was net.

WES HARDAKER: Thanks, Ken. Well, Brad, it wouldn't be meeting if we didn't disagree some of the time, which is just fine with me.

So, a couple of things. First, I want to read one sentence out of the Statement of Work for why this group was formed—and it actually sort of validates everybody's opinion here—which is that "This will allow Internet users to share measured data from their network perspective and help inform Root Server Operators where best to deploy new instances for a better global coverage."

So, a couple of things. One, it does note that a tool is needed in order to measure a local perspective. Right? That's sort of the whole point. And so, one of the things that concerns me, Brad, with your statement was that RIPE Atlas is all that we need. And that has been proven time and again that is not sufficient.

BRAD VERD: No, no, no. If that was the interpretation, I apologize because that's not what the interpretation was meant to be.

WES HARDAKER: Okay. So, that's clearly what I heard. You said, "We just need RIPE Atlas."

BRAD VERD: No, no, no. What I said was the functionality that we have talked about on these calls for a year now, most everybody continues to point to RIPE Atlas in some form or fashion. And if RIPE Atlas was all we needed, we wouldn't be having this discussion.

But what I was trying to point out was that the data in RIPE Atlas which is the data that is—how do I say it—the same context of data, if that makes sense. That that is public. That is public and available to people. So, my question is, why do we want to allow—I mean, not allow. But why do we want to default to an opt-out model where the whole point here is to create a data set that would inform the RSOs better?

WES HARDAKER: Okay. Thank you for the clarification. So, I think we're actually more in agreement than anything else. Yeah. So, RIPE Atlas isn't sufficient, and we can both agreed upon that. The local users might need to wander around multiple networks and run something or other to figure out what they are. Right? It's the same thing of when I run Ping or Traceroute from my local machine. It's because I've discovered something funky.

And I think that you and I also agree that, to a large extent, the type of things that RIPE Atlas can measure are probably all the types of things that we need to measure from our local perspective tool. Whether that is opt-in or opt-out with respect to submission, I think you and I agree there, too, because I, as indicated earlier, suggested a Submit button where it actually becomes manually set to opt in.

And it could be opt-out. I don't know what we want to do in the long perspective there, but I typically am always in favor of a manual opt-in because that's the private-centric way to do this. And I think that we can convince the average person that they want to opt in because it will help make their case that they want more instances near them or something.

KEN RENARD: Thanks, Wes. Fred, did you still have something?

FRED BAKER: Well, yeah. Coming back to Brad's comment that this is primarily for RSOs, I would say that it's also for the folks who would like to host an instance. I can think of a number of cases, primarily at ICANN meetings, where people have come to me and kind of said, "So, how do I get a root server in my favorite location?"

And my question was, first, "Well, do you need one? You go to the neighboring country. Is that a problem?" And so, this tool would ask that question and would necessarily have to be run by the user that wanted to host an instance as opposed to the RSO kind of saying, "Hi. I'm over in Antarctica, and I'd like to figure it out from there."

KEN RENARD: Right. So, that's Use Case 2 and, in general, you could consider that in support of RSOs to make that decision. If a request comes in, an RSO needs to or wants to have this information. So, yes, in theory, a user

could run that tool and share it only with the RSO in question versus publishing to the world.

We're trying to say what this tool should do in a generic sense, not engineering the tool here with respect to the publishing. But we would like to determine a strategy. We have the opt-out or opt-in. Opt-in being strongly encouraging. But I think we all agree that the more data, the more that is published, the better.

Any other comments on the opt-in versus mandatory or optional publishing of the data? Okay. If not, I don't know what the answer is. We may require additional discussion on the mail list of a new proposed text for that.

But for now, let's go on to Wes's comments about the structured and unstructured data, if I can find that. Wes, do you mind talking about that?

WES HARDAKER: I can. I need a 30-second pause. I'm trying to do two things at once, and I'm about to get called in a roll call.

KEN RENARD: Okay. So, while he's doing that, his comment was here in Section 5. So, there are two types of ways that stuff could be stored: structured and unstructured. And we're really talking about the data format. Should this be structured or unstructured?

And my initial thoughts are sort of let the tool implementer decide. But looking at the data that we're collecting, some of it lends itself very well to structured—like the latencies. The routes are at least semi-structured.

Wes.

WES HARDAKER: Thanks, Ken. Sorry for the delay. My point was that if we're going to start collecting data, there's probably a bunch of common things that we need to put into a database: pings, traceroutes, and things like that—like Brad and I were just talking about—as a common format. But then if we want to allow researchers and people to extend this …

If we want to be able to extend this in the future, for example, and add new stuff or to create a system that allows future growth. And that might involve other researchers wanting to do other stuff, too, and submit stuff with permission that has a different data structure or even a text space comment or something like that.

So, I'm trying to figure out how can we make sure that the expandability of the resulting data can continue to [evolve]. And sometimes it is absolutely best to put stuff in a relational data structure where everything is properly typed. And other times, people prefer JSON and MongoDB that actually can change over time. I sort of sit on the fence. I don't like either one.

So, I just wanted people to think about it. I don't necessarily have an answer, but my general thinking is let's define the ones that we know

that we want to collect with the tool. And then you'd just be flexible so that the schema can be altered in the future.

KEN RENARD: Thanks Wes. Ray. Ray, if you're talking, you're talking to the mute.

RAY BELLIS: I'm sorry. I think this absolutely has to be structured, but that doesn't mean it has to be SQL 3rd Normal. It doesn't mean we have to define it here. I think we should leave it as you suggested, Ken, to the implementer to decide what structure is appropriate because, ultimately, the end user's view is [going to be] a JSON style REST API anyway. By which point, the internal implementation of the structure will be invisible anyway.

KEN RENARD: I could see a case where the data comes in at least semi- structured. You've got latency measurements, maybe traceroutes. If there are other things …

Gee, wouldn't it be nice to know that the source of this vantage point is a large university or a small home office? One could take the semi-structured data and push it into a structured repository where, if there's sort of that additional information, it's not necessarily available in ns-table, but it is available in the original text. But we've got to be careful here not to get down to engineering the database itself.

WES HARDAKER: Yeah. In some ways, I apologize for my comment because I think this is a little bit too early in some ways. And this is probably something that should be done by the tool developer, yourself. I think it was you who [said it], Ken.

KEN RENARD: Yeah. I'm just going to write some notes here just saying something along the line of "semi-structured from tool to repo, but it can be presented to users as structured." Okay.

All right. So, the other part of this repository discussion goes down into Section 6, which is Recommendations. We've rewritten that section, and there's … Yeah, thank you, Duane, for your contributions here.

So, we're working with this Recommendation 2 Alternate, and we're recommending that the ICANN Board essentially further development a data repository, to make a specific proposal for a data repository including these 10 things—9.

So, as we read through these A-H looking for comments on any of these—maybe even some additional ones—I think what we're doing here really is more just pointing out what should be addressed during the tool and repository implementation versus any specific solutions.

So, as you read through those, please comment on what we should expand, what we should add or remove from that list. A lot of these are right now implemented as just phrases so that …

I guess we can do some more formalization of the text itself, but a lot of this is just “leave it up to the implementer.” We're not engineering it here.

All right. If there no comments on those, if we want to just go up through the document and look through some of Paul's mostly editorial comments.

Paul, I don't think there's anything here that's really just not acceptable. Just completely, they’re editorial.

PAUL HOFFMAN: The background for this is that I remember at the beginning were concerned with, we didn't want to indicate that someone running this tool would be able to find out what's best or what's needed for them since those determinations are going to be made by the RSOs themselves. So, really, this pass was to take out … Just, some of those slipped in. I mean, it's easy for us to end up saying those things again.

This editorial pass I did was just to remove or soften those so that any determination of the outcome of after I’ve run the tool is purely educational for the person running the tool and not leading towards “need” or “best” or anything like that.

KEN RENARD: Thank you. That's definitely the point we want to get across here. We are informing a decision. We're not making the decision with this tool.

PAUL HOFFMAN: And nor are we assuming that there is a pre-determined "best" at all. Some RSOs might want to use the output of the tool simply to validate their current setup, that is, without saying, "Oh, I care about adding this new one." They may just want to look at/compare their current setup with "What are we seeing from other places around the world?"

So, to try to keep away from any assumption that the tool is going to help with "best" or "good."

KEN RENARD: Thank you. The one comment you have on page 5 here about the texts with text "seems unnecessarily nerdy," I relate to that. And, Paul, with all do love and respect, it's unusual to hear that from you.

PAUL HOFFMAN: I resemble that comment.

KEN RENARD: So, yeah, point taken. That text fit better when we had much more discussion. Okay.

And the next one down, also in Section—I believe—2.1 about "NSID more commonly used …" I think we did discuss this quite a bit before. It seems like hostname.bind versus NSID is the more accepted way to measure.

Does anybody else have any comments on that? Ray.

RAY BELLIS: Yeah. The rationale to stick with hostname.bind is because NSID requires an EDNS option which is potentially more prone to being filtered out. So even though, as far as I recall, every RSO supports both, that was the rationale for picking hostname.bind.

PAUL HOFFMAN: Oh, okay. So, I'll remove that. Verisign, last I heard, actually doesn't support NSID, but Duane keeps saying … I shouldn't speak for him because he's not on the call, but Duane keeps saying, "It's going to come soon. It's going to come soon." But I agree that if there is a filtering out issue and we are concerned about—like, an end user is the one running this—then, yeah, [so this].

NSID is just much easier to describe to someone than CHAOS. But I see your point, so I'll remove that.

KEN RENARD: Okay. And there was a repeat of that somewhere.

PAUL HOFFMAN: I just took it out.

KEN RENARD: Perfect, thank you. And I don't see anything else here that really sticks out as being problematic or that doesn't just make sense.

All right, so I think we've gone through the document. We're still somewhat open with respect to opt-in versus opt-out. We have good

arguments on both sides. At this point I'm uneasy about taking out the last use case about the recursive operator just because we've established that need.

PAUL HOFFMAN: Brad, I have a question for you on that since you were the one who voiced most strongly on it. But I think other people probably agree.

So, I get this testing application and I start to run it. And it says, "Hi. This is for giving you information and sharing that information with the Root Server System." If I don't want to share that information, Brad, would you prefer that the program didn't even run and didn't even give me local measurements? So, that's not opt-out. That's like quit versus opt-out whereas I would prefer that you … Like, I don't want to share it.

BRAD VERD: I don't know. So, look, first I'll say that I'll go along with what the group wants here. I'm just sharing my observation and my thoughts on this and what, as an RSO, I would want. And as an RSO, I want to be as well informed as I can. If that's not what the group agrees on, then so be it.

As far as running the application and, if you're not going to opt in, it just quits on you? I don't know. That seems a bit extreme to me. It seems kind of counter to the "good of the Internet" statement that we made earlier. So, like I said, to me this wouldn't be a showstopper. It's just that I feel like we kind of get stuck on certain topics and we can't get off of them sometimes. But, again, if the group wants to default to opt-out, then I'll go along with group. Those are my thoughts.

PAUL HOFFMAN: So, the reason I was asking is that you brought up a very good point which is once we make this program, it's going to have some sort of user interface and such like that. Opt-out at the beginning is different than opt-in at the end. And so, I think I can support it being opt-out if it's opt-out at the beginning.

That is, "When you run this, you're going to get all of this interesting information and it's going to be shared. But click here if you don't want it to be shared" before they've seen it; as compared to at the end where someone goes, "Ooh, that looks bad. I don't want to share that" because that's exactly … Like, that is the information that you want, Brad, is the not-as-good results.

BRAD VERD: Right.

PAUL HOFFMAN: And I don't know. Andrew, you probably have an opinion here, too, since you were the one who brought it up. But I can easily see something that says "opt out" before running or before collecting data as a reasonable option here.

KEN RENARD: I don't see a reason why, if the user did not want to publish, why they shouldn't actually get the results of the tool. My thoughts based on what Brad had said, even in the last call, was that the user could opt out

of publishing at the point it was ready to publish. If you ran the tool, it collected the data. “If you do not want to publish, click here. Otherwise, it will be published.”

And that's why the way I wrote up the text previous to this meeting was that you had to opt out of publishing. You have the option. You absolutely do not have to publish. But the default if you do nothing is to publish, and with all fair warnings of what the tool does and the fact that it would be published.

The opt-in versus opt-out is really just what defaults if you just press Enter and don't read, which is most mostly what we do with apps for our phones.

Are there any other topics of discussion for this this document as it is? I think the next steps will be to take another pass at the discussion of the repository, mostly with respect to publishing data—yes or no and defaults. Everything else seems to be resolved with the notes that I’ve taken from today. There's no further discussion on the recommendations.

So, as of this meeting, please feel free to continue to comment on the mail list or in the document. I would still like to get this document to a final review within the next month. We're ready to conclude this work party. And I think with that, we can go to Andrew.

ANDREW MCCONACHIE: Thanks. I just have a final AOB before folks hang up. I sent a mail to the RSSAC Caucus list today with a Statement of Work for the RSSAC047

Version 2 Work Party. This is just a plug to please go read it and put some comments in it. That is all. Back to you, Ken.

KEN RENARD: Thanks, Andrew. Is there a scheduled meeting to kick that work party off?

ANDREW MCCONACHIE: No yet. No, we have two weeks to review it and then, depending upon how that review goes, the RSSAC will probably vote to kick that work party off in the next month, I'm guessing. But we'll see how the review goes with the Statement of Work.

KEN RENARD: The work now is just to finish the Statement of Work.

ANDREW MCCONACHIE: Correct.

KEN RENARD: Thanks. All right. With that, I thank everyone for their time and the discussion, and we'll put some new text out there. Appreciate your comments, and that's it for today. Thanks.

OZAN SAHIN: Thank you, Ken. And thanks, everyone, for participating. Tech Support, please stop the recording.

[END OF TRANSCRIPTION]